AUTOMATED DISCOVERY OF FLIGHT TRACK ANOMALIES

Bryan Matthews, SGT Inc. NASA Ames Research Center, Moffett Field, CA
David Nielsen, MCT Inc. NASA Ames Research Center, Moffett Field, CA
John Schade, Kennis Chan, and Mike Kiniry
ATAC Corporation, Santa Clara, CA

Abstract

As new technologies are developed to handle the complexities of the Next Generation Air Transportation System (NextGen), it is increasingly important to address both current and future safety concerns along with the operational, environmental, and efficiency issues within the National Airspace System (NAS). In recent years, the Federal Aviation Administration's (FAA) safety offices have been researching ways to utilize the many safety databases maintained by the FAA, such as those involving flight recorders, radar tracks, weather, and many other high-volume sensors, in order to monitor this unique and complex system. Although a number of current technologies do monitor the frequency of known safety risks in the NAS, very few methods currently exist that are capable of analyzing large data repositories with the purpose of discovering new and previously unmonitored safety risks. While monitoring the frequency of known events in the NAS enables mitigation of already identified problems, a more proactive approach of finding unidentified issues still needs to be addressed. This is especially important in the proactive identification of new, emergent safety issues that may result from the planned introduction of advanced NextGen air traffic management technologies and procedures. Development of an automated tool that continuously evaluates the NAS to discover both events exhibiting flight characteristics indicative of safety-related concerns as well as operational anomalies will heighten the awareness of such situations in the aviation community and serve to increase the overall safety of the NAS. This paper discusses the extension of previous anomaly detection work to identify operationally significant flights within the highly complex airspace encompassing the New York area of operations, focusing on the major airports of Newark International (EWR), LaGuardia International (LGA), and John F. Kennedy International (JFK). In addition, flight traffic in the vicinity of Denver International (DEN) airport/airspace is also investigated to evaluate the impact on operations due to variances in seasonal weather and airport elevation. From our previous research, subject matter experts determined that some of the identified anomalies were significant, but could not reach conclusive findings without additional supportive data. To advance this research further, causal examination using domain experts is continued along with the integration of air traffic control (ATC) voice data to shed much needed insight into resolving which flight characteristic(s) may be impacting an aircraft's unusual profile. Once a flight characteristic is identified, it could be included in a list of potential safety precursors. This paper also describes a process that has been developed and implemented to automatically identify and produce daily reports on flights of interest from the previous day.

Introduction

The NAS continues to evolve as new NextGen technologies and procedures are introduced. A key challenge for aviation stakeholders is to ensure the reaping of the potential benefits of operational efficiencies gained by the new concepts and procedures, while at the same time maintaining the superb track record for NAS safety established over the past several decades. As the NAS changes, it is likely that the safety-related aspects of the system will change as well, and so it is important to actively engage in the discovery of new potential safety risks on an ongoing basis. NASA, in partnership with the FAA, and industry is continuing to develop new technologies and techniques to identify previously undiscovered safety events through the intense data mining of large heterogeneous aviation data sets that are collected on a regular basis. These techniques have the potential to find new safety risks in the system or risks that did not exist previously but are a result of the implementation of NextGen concepts. Combined with more traditional monitoring of safety buffer exceedances and cataloguing of known safety-related incidents, this approach helps to provide a more holistic view into the safety of the NAS.

This paper presents the next step in the development of advanced data mining algorithms as applied to high fidelity surveillance and trajectory data. It builds upon previous research [1] by adding more features to the mathematical models to discover previously unknown safety-related events. This research expands into operations in the New York Metro area and at Denver International Airport to see how the discovery algorithms can function in both highly congested and dynamic weather impacted environments. As with previous efforts, subject matter expertise is incorporated into the research to provide specific domain knowledge of operational procedures and to understand the safety implications of the discovered events. In addition, the evaluation of algorithm results is made clearer by the inclusion of (ATC) voice communications from frequencies involved at the time of the event occurrence.

This paper is organized as follows: First we present an overview of the Performance Data Analysis and Reporting System (PDARS) which delivers several capabilities to enhance this research including serving as the source of the trajectory information. Next we discuss the state of the previous research efforts in this area. We then present the primary algorithms used for the discovery of safety events and the specific data used for the research. Since the data processing and handling is quite involved, we give an overview of the end-to-end system for algorithm application and introduce a prototype for incorporating these techniques on a daily basis. We then present the primary results of the research including 11 actual traffic scenarios that were identified as operationally significant anomalies along with a brief safety analysis for each. Although 5 out of the 11 scenarios involve go-arounds, which are typically recorded in the control tower's daily logs and captured by daily PDARS reports, the algorithm and the subject matter expert's reviews also provide detailed insight into what factors contributed to the anomalous event. Finally we discuss the conclusions and introduce ideas for future research goals.

Background

PDARS Program

For over a decade, the Performance Data Analysis and Reporting System (PDARS) has continued to provide FAA organizational managers and decision makers with "actionable" information regarding the efficiency and safety of the NAS. PDARS is a product of collaborative research between NASA and the FAA that was recognized for its excellence by receiving the NASA Administrator's Turning Goals into Reality (TGIR) award in 2003 and achieving full technology transfer from NASA to the FAA in 2005[2]. The PDARS program is managed by the FAA's Air Traffic Organization Office of System Operations Services and is heavily used operationally by over a dozen organizational units within the FAA, many on a daily basis.

PDARS consists of an ever-evolving data collection, processing, reporting, and dissemination platform able to accept nearly any surveillance or positional data and merge that with other geo-referenced or contextual aviation-related data (e.g. weather, terrain, or schedules). The system routinely produces analysis products including reports and visualizations that provide detailed operational insight to decision makers at virtually any level in a complex Air Navigation Service Provider (ANSP) organization such as the FAA. The development of PDARS has been from the beginning and continues to be driven by the needs of the user base: those actively involved in direct operation of the NAS and the associated challenging areas such as safety, efficiency, and environmental concerns [2]. PDARS' flexible reporting structure produces over 1,500 reports daily, many of them safety related such as go-arounds, Class B airspace excursions, interacting runway operations, and turns-to-final.

Key PDARS capabilities used in this research are its routine collection and processing of large surveillance-based trajectory information sets, categorization of key flight parameters such as runway utilization, its ability to compute additional geospatial measures on aviation data sets on a large scale, and its suitability as a technology transfer platform for new technologies. In particular, the development of a prototype daily anomaly report is one important outcome of this research. The data mining algorithms presented in this paper make perfect candidates for incorporation into one or more PDARS "anomaly" reports, which could be produced on a daily basis.

Current State-Of-The-Art Safety Monitoring

Many of the existing safety monitoring technologies used today in the NAS are based on the ability to define the characteristics of high-risk safety events and utilize current sensor measurements to detect safety risks. Safety tools used by the FAA that monitor continuous loss of standard separation such as the Operational Error Detection Program (OEDP) [4], which look at loss of separation en route, have been in place for more than a decade. More recently deployed is the Traffic Analysis and Review Program (TARP), which monitors Terminal Radar Approach Control Facilities (TRACON) operations and provides the Comprehensive Electronic Data Analysis and Reporting (CEDAR) tool, the way to clearly define when a safety incident occurs and to what degree it was unsafe. A report filed by the Office of Inspector General in February 2013 [5] discusses the effectiveness of TARP and CEDAR; however, it notes that there is a shortcoming in the tool's ability to completely capture all loss of separation incidents as compared to what was reported in the voluntary Air Traffic Safety Action Program (ATSAP). The report also states that the FAA was unable to fully review all cases to determine if a valid loss of separation had occurred. These limitations can lead to a significant discrepancy between the safety risks that are monitored and those that are actually happening in the airspace. Since these tools are designed to identify specific safety risks, they can be very effective in detecting the known safety events being monitored; however, they are not looking for events outside of their predefined scope. Airline operated Flight Operational Quality Assurance (FOQA) programs take a similar approach to detecting and tracking safety events within an airline's fleet. By using predefined safety events, rule-based exceedance checks are used to determine to what degree certain aspects of a flight were safe or not and what actions can be taken, such as enhanced training or internally circulated newsletters, to mitigate these known risks.

While using the predefinition of safety events methodology is very effective in monitoring safety incidents and their frequency, this approach does not have the flexibility to detect new or unknown safety issues. Working groups such as the Commercial Aviation Safety Team (CAST)[1] work with industry and government agencies to address and reduce newly identified risks in the NAS. These risks may be identified by pilot safety reports and analysis of the flight data and/or accident/incident investigations; however, this typically is a manual process. Over the last decade, NASA and other institutions such as MIT have invested in developing data mining techniques to discover safety events in FOQA data [3],[6],[7],[8],[9]. Some of these algorithms have been evaluated on large domestic FOQA data repositories by MITRE through the Aviation Safety Information Analysis and Sharing System. Although FOQA data contains detailed measurements from a fleet of aircraft and can be used to address many safety questions, one drawback, besides only being collected from participating carriers, is that it is limited to single aircraft operations and does not fully capture the interaction between multiple flights in the airspace at the same time. Recently, these algorithmic concepts have been adapted to the radar track data [1] where safety events that involve safely spacing and sequencing of flights were detected and validated by subject matter experts from the Southern California TRACON (SCT). The work discussed in this paper builds upon the previous SCT radar track anomaly detection study, while incorporating new features that help to characterize the interaction between multiple aircraft. The tool is also validated by data from two additional TRACON facilities covering four major domestic airports. This paper also discusses how the algorithm has evolved from a research tool into an integrated prototype daily anomaly reporting capability.

Multiple Kernel Anomaly Detection

The algorithm selected for this study, Multiple Kernel Anomaly Detection (MKAD), uses a one class Support Vector Machine [1] architecture, which is used extensively for anomaly detection in the field of data mining. The method learns a decision boundary separating nominal and anomalous data points based on a pairwise kernel (similarity) matrix and can identify statistically significantly anomalous examples. One of the key components to SVMs is the

[1] http://www.cast-safety.org/

kernel function that defines the similarity between two vectors. Choosing a measure that poorly distinguishes between these vectors will significantly impede the performance of the algorithm. In our previous SCT study [1] we used the cosine similarity as the kernel function between two trajectories because it possessed a straightforward geometric interpretation between the flight trajectories and did not have any hyper-parameters that needed to be optimized. For this study we replaced the cosine similarity kernel with the radial basis function (rbf) (Eq. 1), which has been very popular in the machine learning community for many years.

Kernel Function: $\kappa_n(\vec{x}_i, \vec{y}_j) = exp^{-\frac{\|\vec{x}_i - \vec{y}_j\|}{2\sigma^2}}$ (1)

Though, the cosine similarity kernel could still be used for all features, the rbf kernel function was chosen because an additional feature (the distance to the nearest aircraft vector) had been added to the feature space. With this additional feature the geometric interpretation of the cosine similarity function that pertained to the previous trajectories was no longer appropriate. On the other hand, the rbf kernel can be abstracted to any multidimensional vector space. Even though MKAD has the ability to combine kernels from different kernel functions, in practice it has been observed that keeping the kernels the same whenever possible helps to maintain similar distributions across the kernels and helps produce anomaly feature contributions with more consistency. In this case the rbf kernel was appropriate for all features. One additional complexity that the rbf kernel has over the cosine similarity kernel is the hyperparameter σ that determines the width of the Gaussian distribution over the vector space. If σ is too large then all vectors appear similar and if σ is too small then all vectors appear dissimilar. The approach that was used to tune σ in this study was to compute a 1000 x 1000 kernel from a random sample of vectors that had been z-scored (zero mean and unit standard deviation). The kernel was recomputed for the same set of vectors over a range of σ values to determine the σ that yields the kernel distribution with the minimum variance, which therefore corresponds to a wider spread of similarities across the kernel space. This wide spread distribution tends to yield good distinguishing ability and is used to identify the relative order of magnitude for σ. Cross validation is typically done in data mining to ensure the σ is optimal, however for this study, labeled examples of anomalous flights are not known beforehand and therefore cross validation was not an option. After performing this search 5-folds (where data is randomly subsampled for each iteration), a best σ was computed by taking the mean σ value across the 5-folds. This process was observed to be stable with tight bounds across the runs and has worked well in other best σ searches for other data mining applications.

Once the best σ values are calculated the kernels for each trajectory (latitude, longitude, altitude, and distance to nearest aircraft) for similar runway destinations are computed from the z-score normalized training data. The novel aspect to MKAD is in its ability to combine the features in the kernel space. This is simply done by taking the weighted average of each kernel shown here in Eq. 2, where W_m represents the weights of the m-th kernel and κ_m is the m-th kernel function. For this study all kernels were given equal weights of 0.25.

Combined Kernels:

$$\kappa(\vec{x}_i, \vec{y}_j) = \sum_{m=1}^{n} W_m * \kappa_m(\vec{x}_i, \vec{y}_j) \quad (2)$$

After the kernel is combined the 1-class SVM can be solved given the following optimization problem and constraints shown in Eqs. 3 and 4.

Minimize:

$$Q = \frac{1}{2}\sum_{i,j}^{\eta} \alpha_i \kappa(\vec{x}_i, \vec{y}_j)\alpha_j \quad (3)$$

Subject to:

$$0 \le \alpha_i \le \frac{1}{\eta\upsilon}, \sum_i \alpha_i = 1, 0 \le \upsilon \le 1 \quad (4)$$

Where η is the number of trajectories in the training set and υ is provided by the user and corresponds to the maximum fraction of data assumed to be anomalous (in this study υ was 15%). After solving for α, given the constraints, the non-zero α values are considered to be the support vectors and define the hyperplane used to separate the anomalous trajectories from the normal. To determine the distance the hyperplane lies from the origin, ρ needs to be calculated (shown in Eq. 5).

$$\rho = \frac{1}{length(\alpha)} \sum_{j \in \alpha > 0} \sum_{i \in \alpha > 0} \alpha_i \kappa(\vec{x}_i, \vec{y}_j); \qquad (5)$$

$$where: \rho \geq 0$$

The anomalous examples are rank ordered by the algorithm based on their distance to the hyperplane. The anomalies are located on the negative side of the hyperplane, whereas the nominal examples are on the positive side. The severity of the anomaly is determined by its distance from the hyperplane. Only negative examples are marked as anomalous. The formula for calculating the rank order scores is shown in Eq. 6.

$$Score(y_i) = \sum_{i \in \alpha} \alpha_i \kappa(\vec{x}_i, \vec{y}_j) - \rho \qquad (6)$$

Once the models for each runway are learned on training data from a sliding window of previous *N* days, the models are applied to a test day to determine which flights are anomalous within that day. After the flights are identified the contributions from each kernel can be linearly computed and are used to determine the factors that contributed to the flight being labeled as anomalous (for example flights with high proximity to neighboring flights may have high contributions from the distance to nearest aircraft vector).

Data Management Process and Prototype

The surveillance data used in this study comes from four ATC facilities. The New York Terminal Radar Approach Control (N90) and the New York Air Route Traffic Control Center (ZNY) provide surveillance data for the New York area. The Denver Terminal Radar Approach Control (D01) and the Denver Air Route Traffic Control Center (ZDV) provide surveillance data for the Denver area. With the FAA approval, NASA was given access to PDARS data from the 2013 calendar year. Approximately 386,000 flights were analyzed to obtain the final results listed in the results section below.

Figure 1 illustrates the data processing flow from data collection through merging, filtering and anomaly detection. Starting at the top center of the figure, raw ATC data collected from the ATC facilities are first processed to create flight trajectories. Some surveillance data from the FAA facilities contain data from multiple radar sensors. In this case, N90 and D01 are both multiple radar systems, which contain data from 5 and 10 radar sensors respectively. The multiple sensors systems create additional complexity when producing high-quality flight trajectories for analysis as the coverage for those sensors can overlap. During the processing step, the PDARS system selects the best radar hits to use based on many different criteria in order to produce the best quality of four-dimensional (latitude, longitude, altitude, and time) trajectories for flights. The resultant data provides analysis ready trajectories between Air Route Traffic Control Center airspace boundary and the Terminal Radar Approach Control boundary. Since this study is focused on finding unusual patterns in commercial Instrument Flight Rules (IFR) aircraft, Visual Flight Rules (VFR) flights with beacon codes from 1200 to 1299 and military flights are removed from the data for analysis however, they are considered when calculating the distance to nearby aircraft. The additional benefit of removing those flights is that military and non-discrete code VFR flights typically will have unusual flights paths as compared to commercial flights, and by removing those flights the tool is expected to yield more relevant results during the data mining process. Also during the processing step, additional flight information such as destination airport and landing runway are computed at the same time.

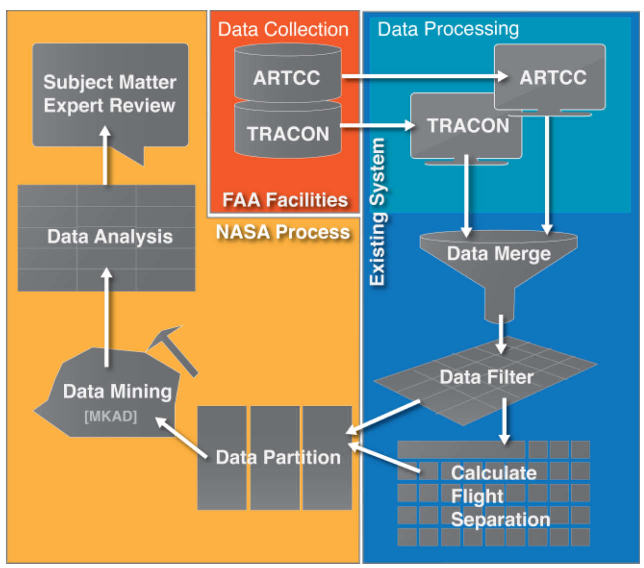

Figure 1 Data Processing Flow

The flight trajectories are then used to calculate the minimum separation between flights. The delta vertical and lateral distance between track points for each pair of flights with the same timestamp are compared. A daily time-series report showing where each of the two flights is located during the time overlap is generated. This daily report is then filtered to remove all entries with vertical separation over 2000 feet and lateral separation greater than 6 nautical miles. Subject matter experts selected these maximum values to retain operationally significant separation information to capture trajectories indicating how rapidly aircraft were converging under the upper bounds, while reducing the files to a manageable/searchable size. For each trajectory, from 30 nautical miles (NM) out from the destination airport, the minimum separation is found and used to create four-dimensional trajectories: latitude, longitude, altitude and distance to nearest neighboring flight. If no separation values are within the thresholds, the threshold values are used. These four features are then averaged over half nautical mile intervals from 30 NM to the runway threshold based on distance traveled and are partitioned into sets of trajectories landing at each airport on each day. This results in having uniform trajectories with fixed vector lengths because of the half-mile binning and the fixed 30 NM distance traveled.

For every day and airport runway, the previous days are compared against each other to create a model of nominal behavior for each destination runway and used by the MKAD algorithm. Runway usage (due to hardware memory constraints) limits the number of preceding days that can be considered; for this study we used the following training sets: DEN (120 days), EWR (30 days), JFK (60 days), and LGA (50 days). The algorithm builds a training model based on the trajectories from the previous rolling window and tests on the day of interest to compute the anomalies. The anomalies are reported to data analysts for examination. Trajectories of interest are investigated further utilizing the Graphical Airspace Design Environment (GRADE), a graphical analysis tool. The overall traffic flow is visualized with the tool to obtain a better understanding of the airspace and traffic flows for each situation. After some flights of interest have been identified by the analysts, the voice data, when available from LiveATC.net, is analyzed to find the relevant communications that pertain to the flights of interest. Transcripts of the recordings are done by hand and used to provide context for the most significant scenarios. This information is used to further understand the scenarios and the key characteristics are summarized with animations and presented to the subject matter experts familiar with the airspace.

The refined algorithmic tool from this study is deployed as a prototype into a lab demonstration of the PDARS environment to demonstrate the feasibility of transferring this technology from a research environment into an operational environment by automatically generating a daily report of the algorithm output. The prototype daily anomaly reporting program follows a similar flow but has been automated to run when the trajectory file with the previous day's flights is available on the staging server. Additionally flights windowed around the same time frame from the previous year that have already been processed are added to the training corpus along with the previous month's flights to help build a more comprehensive model and to preserve seasonal effects in the training data. For the first case study the report prototype is customized to analyze flights landing at DEN. Once subject matter experts have evaluated the results using the criteria used in the results section of this paper, additional airports can easily be set up to be monitored.

Results

Approximately 90 flights were given to subject matter experts for further analysis after visually inspecting from the list of anomalies generated by the algorithmic tool. Out of these, 33 were deemed to hold some operational significance. The remaining 57 were considered statistically significant anomalies but were not considered by the subject matter experts to pose a significant safety risk. The operationally significant flights were presented to local controllers and subject matter experts from D01 and N90 familiar with the everyday operations at the respective facilities. For brevity, this section will cover 11 representative scenarios of the 33. Each scenario discussed below provides a short description of the operational situation, offers the subject matter experts'/controllers' feedback for the possible explanation(s) of what may have led up to the flight's unusual behavior and provides a description of each scenario's potential relevance to safety risk. The scenarios were reviewed by FAA Safety personnel and checked against TARP and CEDAR logs. The expectation was that if a TARP event indicated a loss of separation it would create an Electronic Occurrence Report (EOR). In order to reduce unnecessary or nuisance alerts, facilities have the ability to create rules that define exemptions near the airport for alert reporting. The TARP/CEDAR review board then makes a determination as to whether the exemption may be approved. Flights that qualify for an exemption or are single events will not generate an EOR. The PDARS go-around reports were also checked to determine if a flight was caught by the current automatic reports.

In the following descriptions the aircraft identified with the unusual trajectory is referred to as TGT AC (i.e., the "target aircraft"). Other flights in the airspace will be referred to as FLTXX. The Extended Runway Center (ERC) lines (where appropriate) are shown in a GRADE-based graphic that illustrates the scenario to give a sense of horizontal alignment with the runways for the aircraft's turn to final. Pertinent information contained in aviation routine weather reports also known as METAR weather reports for the hour of an aircraft's arrival at the destination airport are provided for scenarios when relevant.

LGA Scenario 1

Description: A Bombardier CRJ700 series (TGT AC) flight landing on runway 22 at LGA is given an extended right downwind leg for sequencing to follow an arrival on a straight-in approach to the same runway. TGT AC intercepts runway 22 ERC at 10 NM from runway threshold and is 38 knots faster than the proceeding aircraft and in trail by 3 NM. Over the next 2.5 minutes, TGT AC is able to decelerate to a speed closer to the preceding aircraft; however by this time in trail distance has decreased to 1.5 NM. LGA tower voice recordings reveal that TGT AC is issued a go-around by the LGA tower controller due to the overtake situation. Figure 2 shows the sequential closure of in-trail distance and speed differential between the TGT AC and FLT01 that indicate an impending overtake on a single runway operation. The situation was mitigated by the controller executing a go-around.

Explanation: TGT AC is on a vector for an Instrument Landing System (ILS) approach and the traffic pattern is extended out to approximately 14 NM to allow merging of flights on straight-in to runway 22. TGT AC is considerably faster than FLT01 and the TGT AC may have kept its speed up to enable an earlier traffic merging and/or to avoid any additional delay. The voice recordings indicate that the LGA tower controller instructed the TGT AC to go-around with a clearance to maintain runway heading and climbing to 2,000 feet. If wake-turbulence were an issue the LGA tower controller would most likely have initiated a go-around earlier. The tool flagged this anomaly due to the unusual proximity to the nearest aircraft as the top contributing factor at this point in the flight.

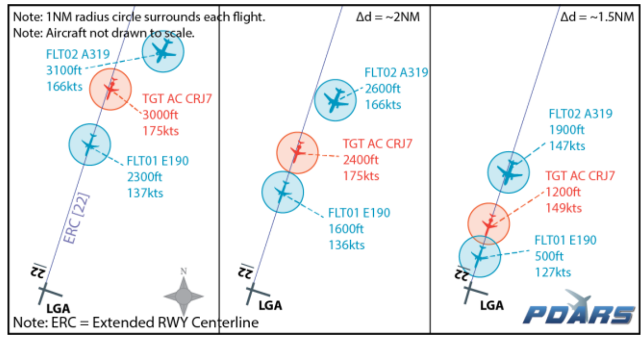

Figure 2 Impending Overtake Scenario.

Safety Review: Apparently the trailing aircraft was not able to reduce airspeed in time and/or fly the aircraft to maintain sufficient spacing (this could have been due to the aircraft not being configured yet). The controller resolved the situation by issuing a go-around and re-sequencing the TGT AC for another approach. CEDAR identified an EOR for the TGT AC. There was a manually written report associated with the TGT AC stating a turbo-jet go-around within a 0.5 NM of the runway. PDARS also captured this flight in the daily LGA go-around report.

LGA Scenario 2

Description: An Embraer E-170 aircraft (TGT AC) descending from 12,000 feet on a vector for a landing on runway 22 at LGA stops descent at 5,000 feet when approached from the opposite direction by a VFR flight Piper PA-28 Cherokee (FLT01) at 4,500 feet and 1.5 NM southwest. TGT AC then immediately executes climb back to 5,500 feet to establish needed separation.

Explanation: The LGA Sector controller receives TGT AC descending en route to LGA. At the same time, a PA28 aircraft is approaching from the opposite direction over Long Island Sound and descending on a VFR assigned beacon code. TGT AC is observed descending to 5,000 feet and then climbing back to 5,500 feet. No voice recordings were available for this event and it is unclear if the LGA controller issued climb or if the aircraft received a Traffic Collision Avoidance System Resolution Advisory.

Figure 3 illustrates the situation in which aircraft are separated by less than 1,000 feet and 3 NM while traversing in opposite directions within terminal airspace. This situation is corrected when TGT AC climbs back to higher altitude thus re-establishing needed separation criteria. The tool flagged this anomaly due to the unusual proximity of TGT AC to a nearby aircraft as the top contributing factor at this point in the flight.

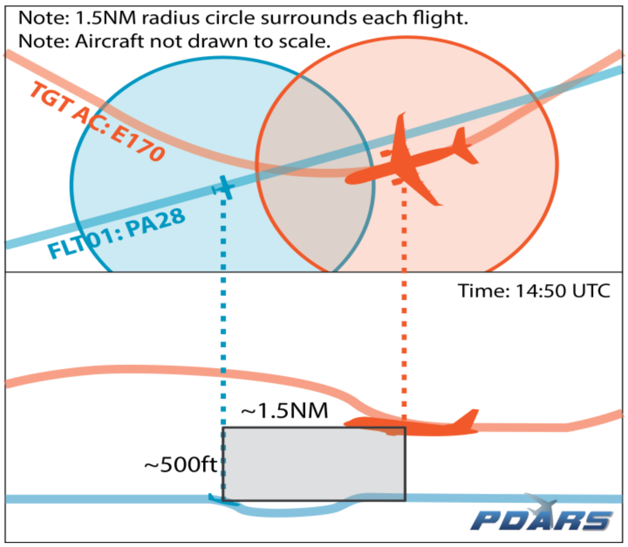

Figure 3 Loss of Separation Scenario.

Safety Review: The airspace stratification is very complex in the area where the aircraft are transitioning, with aircraft being worked by multiple controllers. Our developed data mining tool could be useful in identifying high-risk IFR/VFR crossing "hotspots". A request was made to the CEDAR logs, but an EOR was unavailable for this event.

JFK Scenario 3

Description: A Boeing 777-200 (TGT AC), an overseas arrival to JFK, executes an approach for landing runway 31R. The TGT AC is at a very high ground speed and high angle of intercept as the aircraft approaches the 31R ERC outside the Outer Marker (OM). At the time of approach, the winds were from the north at 16 gusting to 22 knots, which could have contributed to the aircraft's excessive ground speed. This high-energy approach results in the TGT AC overshooting the final approach course by more than 5,000 feet and when an attempt to make a visual approach to runway 31R is unsatisfactory, the pilot initiates a go-around.

Explanation: The aircraft may not have been configured to land and makes the decision to go-around. It is also feasible that the TGT AC may have lined up incorrectly for the wrong runway (i.e., 31L). Available voice recordings only indicate pilot stating "going around" at the time of the maneuver. The aircraft's late turn to intercept ILS 31R with ground speed at 252 knots is 70 knots greater than the other aircraft observed executing the same approach procedure during this time period. On the second approach the aircraft is told by the controller to expect ILS 31R and the controller queried if the pilot had airport in sight. The pilot requests ILS 31R and conducts a late turn and overshoots ERC again before landing. Figure 4 shows the aircraft intercepting the final approach course at too high of a ground speed and at an excessive angle of intercept. The outcome is an overshoot of the ERC and the inability to maneuver the aircraft for a safe landing results in a pilot initiated go-around. The tool flagged this anomaly due to the unusual altitude profile as well as contributions from the latitude shift as the top contributing factors.

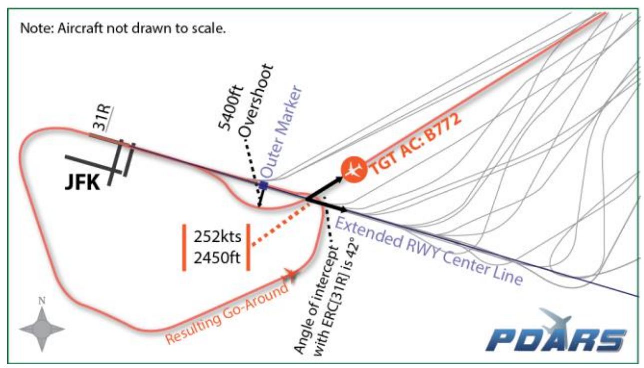

Figure 4 High Speed and Angle of Intercept Scenario.

Safety Review: Desirable angle of intercept and speed of intercept of the ERC are 30 degrees or less and 180 knots or less, respectively. Unfamiliarity with airport, existing wind conditions, pilot technique, long haul fatigue, and/or failure to configure aircraft for landing in time may have been factor(s) in this occurrence. This flight was captured in the daily PDARS JFK go-around report.

JFK Scenario 4

Description: An Airbus A320 (TGT AC) overshoots the ERC at JFK during a midnight operation for landing runway 04L. The final approach is initiated 10 NM from the runway threshold at an excessive intercept angle, ground speed, and altitude as compared to other flights observed landing that day. Vertical separation is maintained from an adjacent heavy aircraft already established on ILS 04R. A controller voice recording indicates TGT AC is instructed to maintain 2,500 feet until established on centerline of 04L before descending to complete visual approach.

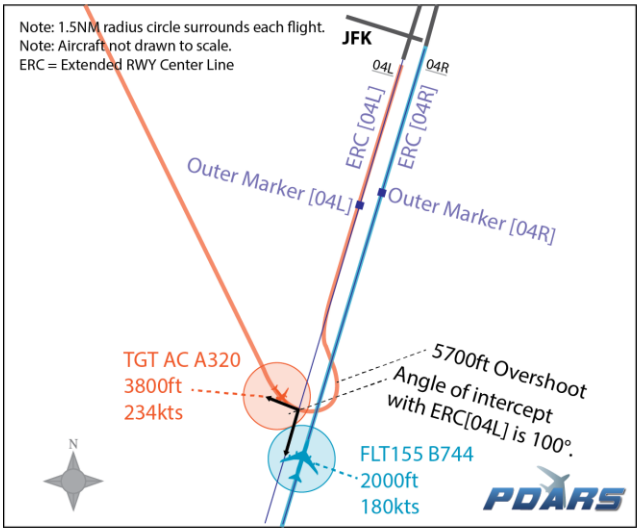

Figure 5 Visual Approach with ERC Overshoot Scenario.

Explanation: TGT AC is on a vector for a visual approach[2] to runway 04L at the same time as a heavy arrival (FLT155) is executing an ILS 04R approach on adjacent parallel runway. The ground speed (234 knots), altitude (3,800 feet) and angle of intercept (greater than 100 degrees) are all contributors to the overshoot of the target ERC. The voice recording indicates that the controller apprised TGT AC of the FLT155 and that the TGT AC confirmed visual contact. The JFK controller separated aircraft vertically, but may have delayed turn for operational concerns thus issuing clearance to intercept localizer too late. Figure 5 illustrates excessive ground speed, altitude, and angle of intercept by a late night JFK arrival conducting a visual approach. The result was an overshoot of the desired ERC with vertical separation maintained from aircraft on the adjacent parallel runway. The tool flagged this anomaly due to the proximity of TGT AC to a nearby aircraft and the unusual southern latitude shift as the top contributing factors at this point in the flight.

Safety Review: The TGT AC was cleared for a visual approach to runway 04L. Excessive ground speed, altitude, and angle of intercept all contributed to the overshoot as well as potential late controller action. Late night operations could have affected the pilot's performance in attempting a visual approach to the landing runway. A CEDAR log request was made, but an EOR was unavailable for this event.

JFK Scenario 5

Description: In this scenario, a routine daily short-haul cargo flight, a Beechcraft 1900 twin-turboprop (B190) from Canada to JFK that normally arrives between 11:00 PM and 01:00 AM is examined. The initial flight plan for the daily B190 flight is a routing to a fix 60 NM north of JFK for sequencing to primary arrival runway 22L at JFK. Generally when operational circumstances allow, the original flight plan is either amended or cancelled and a direct routing for VOR or visual approach to 13L is requested to shorten flight time and expedite taxi time to the cargo terminal (located at the north end of JFK). Upon approaching JFK for landing 13L, airport conditions may not allow an immediate landing on desired runway and various flight maneuvers (see Figure 6) are used to redirect the aircraft's approach to another runway that still provides a landing roll-out as near to cargo terminal as feasible. Weather could have been a factor for the flight that was re-sequenced from 13L to 22L, which required a close in 270 degree turn. Weather in this instance was 4 SM visibility and light rain as reported in the METAR data.

Explanation: In addition to switching to various runways, flights can be required to make 360 turns and S-turns to accomplish the landing setup to amended runway assignments. These arrivals utilize close-in rapid descents from high altitudes when approaching JFK for 13L and characteristically remain for the most part above standard runway glide slope. As a cargo flight, passenger comfort is not a concern and may be a factor in allowing these types of approaches. The tool flagged this anomaly due to the

---

[2]Section 5-4-23 Visual Approach. (n.d.). In *the Federal Aviation Administration's Aeronautical Information Manual* (July 24, 2014). Retrieved from http://www.faa.gov/air_traffic/publications/ATpubs/AIM/aim0504.html#57

unusual sharp drop in altitude as the top contributing factor in all of these flights.

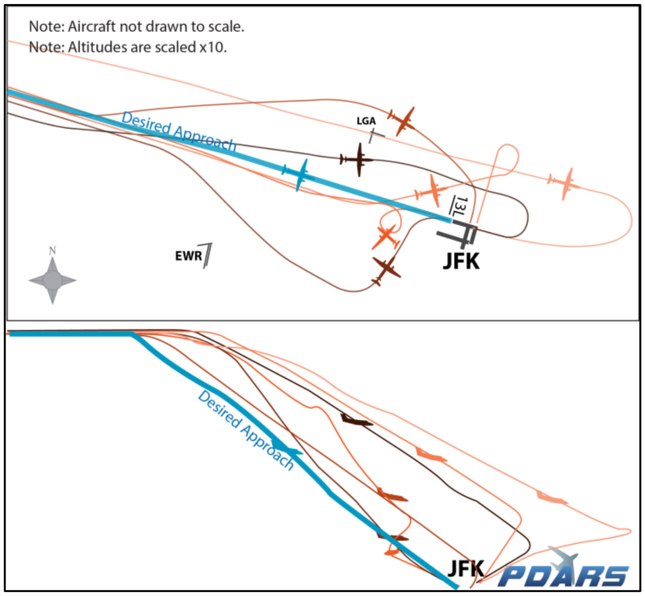

Figure 6 Unusual Approaches for Same Cargo Flight Over 7 Different Days.

Safety Review: These flights do not approach the airport via the conventional approach used by the majority of arrivals and in addition to making a close-in high-energy descent, a wide spectrum of maneuvers are utilized to re-route the aircraft to other runways. These maneuvers include: path stretching, 360 turns, and S-turns all conducted at low altitude and late at night. Since this routing does not appear to be normal or standard operating procedures (SOP), not all controllers may be aware of how best to handle this nightly arrival. Using the data mining tool to identify unusual behavior could help facilities modify their SOPs so all controllers have an approved, coordinated, and documented procedure to handle these frequent unusual routings.

EWR Scenario 6

Description: A Beechcraft Super King Air 200 (TGT AC) following an Embraer ERJ 145 (FLT899) for landing on runway 22L at EWR. Voice recordings indicate that the pilot was unable to ascertain that a larger landing aircraft is in sight although the in-trail distance from the landing flight is 3.6 NM. There is a 20-knot tail wind during the approach and METAR reports that overcast conditions exist which could contribute to pilots' inability to see the preceding aircraft. The EWR controller issues the TGT AC a go-around to avoid a potential wake turbulence situation under the existing approach circumstances. The TGT AC executes go-around and is re-sequenced for subsequent landing on same runway. Figure 7 shows an aircraft making a go-around initiated by the EWR controller as TGT AC reported landing aircraft not in sight.

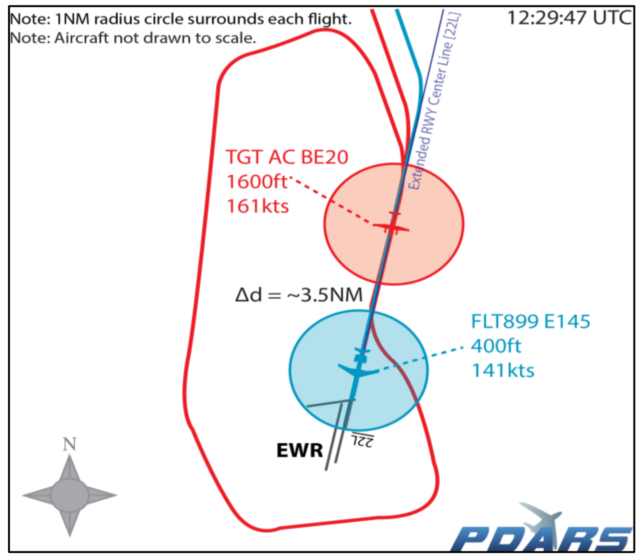

Figure 7 EWR Controller Initiated Go-Around Due to Lack of Visual Contact Scenario.

Explanation: Approach turn-on angle and altitude at ERC intercept are both consistent with other approaches that day, but an overtake situation is progressing as the TGT AC is 20 knots faster than FLT899. TGT AC (a small wake vortex category of aircraft), following FLT899 (a large wake vortex category aircraft) is unable to maintain the required wake turbulence criteria as well as not having FLT899 in-sight. The EWR tower controller most likely initiates the go-around due to wake turbulence considerations and because the TGT AC is unable to visually identify FLT899 on final. The tool flagged this anomaly due to the unusual longitude trajectory during the go-around.

Safety Review: The trailing aircraft is not able to establish visual contact on the landing aircraft and the wake turbulence considerations become an issue. The controller resolves the situation by issuing a go-around and re-sequencing the TGT AC for landing. This aircraft was not included in the PDARS daily go-around report at EWR because the report filters out most general aviation aircraft.

EWR Scenario 7

Description: This fix is normally used to transition the aircraft for a right downwind approach to EWR runway 22L. In this case, TGT AC is S-turned to establish a left downwind for EWR runway 22L. This maneuver required the TGT AC to tunnel under EWR departure just departing from 22R as well as requiring FLT1623 to stop its left turn until sufficiently clear of the TGT AC. Figure 8 shows the re-routing of an EWR arrival that normally would utilize a right downwind for landing. In this case, the flight is rerouted to the east side of EWR for landing, front-loading a heavy arrival rush by delaying inbound flights from the north and west.

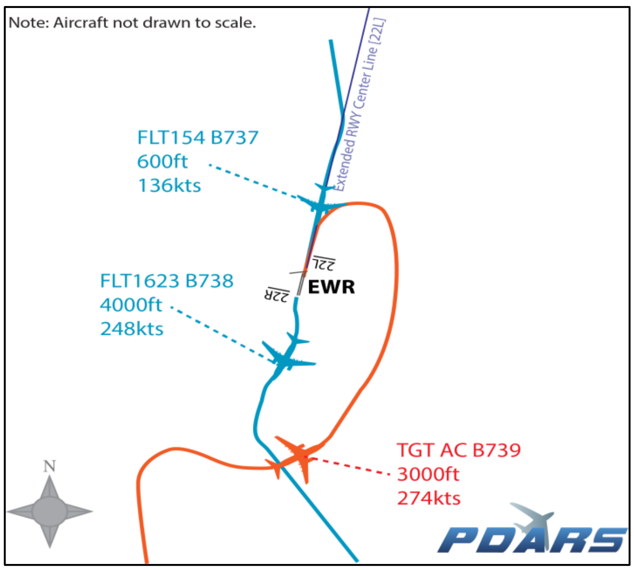

Figure 8 Unusual Re-Routing Arrival Event.

Explanation: The TGT AC is a San Juan departure filed over a waypoint 80 NM southeast of EWR. Other San Juan departures that day on a similar routing were also turned early but the majority assigned a right downwind for Runway 22L. Just after the TGT AC landed a heavy approach stream from the north was experienced followed by a heavier arrival rush from the west where flights were both path stretched to the north or issued holding prior to release for landing. There were no voice recordings available for this event. The tool flagged the unusual latitude trajectory as the top contributing factor.

Safety Review: The left downwind to EWR 22L was used to possibly expedite and ease the sequencing for 22L prior to the heavy arrival rush forthcoming from north and west directions. The left downwind maneuver required crossing under EWR departure and to operate in the vicinity of two arrival flights inbound for landing at both EWR and TEB. This maneuver requires a rapid descent and abrupt turn in order to descend the aircraft below the altitude of the 22R departures. The potential for a conflict with a departure is increased.

EWR Scenario 8

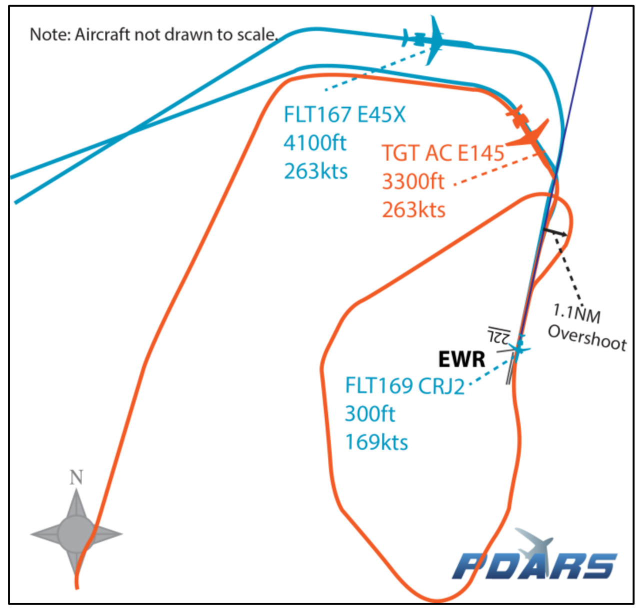

Figure 9 Go-Around with ERC Overshoot Scenario.

Description: An Embraer E145 (TGT AC) arriving from a westerly direction experiences a slight overshoot (300 feet) after crossing EWR 22L ERC before being successfully established on the 22L localizer. The TGT AC's approach to the ERC (3 NM outside the OM) is at an altitude (2,900 feet), an intercept angle (40 degrees) and a ground speed (257 knots) consistent with other observed flights on the same flight trajectory that day. These other flights were able to execute a landing at EWR, whereas the TGT AC initiated a go-around upon reaching the threshold. The final approach for the TGT AC exhibits flight characteristics very similar to other flights landing at EWR that day. Figure 9 shows a go-around of the TGT AC after approaching the runway threshold at an altitude and ground speed consistent with previous EWR arrivals.

Explanation: The distance to the preceding aircraft FLT169 is 8 NM, so wake-turbulence would not be a problem. The tower voice recordings reveal that the pilot initiated the go-around. Upon the second approach the tower controller asks if wind shear was the cause; however the pilot responds by stating that the aircraft was "just too high". The tool flagged this anomaly due to the unusual latitude trajectory as the top contributing factor during the go-around.

Safety Review: A go-around was executed successfully at the runway threshold and on the second approach the aircraft experienced an overshoot of 1.1 NM due to an excessive ERC intercept angle. Subsequently the TGT AC was not established on the localizer until 1.5 NM inside of the OM. Aircraft altitude at the runway threshold and ground speed were nearly identical to the first approach. This flight was captured in the daily PDARS go-around report for EWR.

EWR Scenario 9

Description: A Boeing 737 (TGT AC) conducts a visual approach to Runway 11 at EWR during afternoon converging runway operations and "ties" with an A320 (FLT1507) approaching Runway 22L from the north. The tower controller advises the TGT AC to execute a go-around due to conflict with arrival of FLT1507 on 22L. The TGT AC is re-sequenced with later EWR arrivals inbound to 22L.

Explanation: The TGT AC may have been vectored and issued a late turn to final resulting in a tie with FLT1507. Although a passive visualization tool Converging Runway Display Aid (CRDA), which can be used to assist tower controllers in the spacing between Runway 11/22L converging approaches, was in use (according to the tower logs), it may not have been utilized properly making the coordination and timing of the two approaches more difficult. Voice recordings revealed that the EWR tower controller advised the TGT AC of conflict with FLT507 and sent the TGT AC around. Figure 10 shows a go-around of the TGT AC after criteria for continuation of converging approaches to Runways 11/22L at EWR were not met and the tower controller subsequently had TGT AC execute a go-around. The tool flagged this anomaly due to the unusual latitude trajectory as the top contributing factor during the go-around.

Safety Review: Use of converging approaches requires that adequate spacing be available between the two aircraft executing the simultaneous approaches. In this case, the tower controller determined the necessary criteria were not being met for this EWR operation and directs the go-around for a re-sequencing of the TGT AC. This incident was not an isolated event while the CRDA was in use; another flight on a different day was re-sequenced to the crossing runway due to a similar potential "tie" situation. There was no indication in the tower logs that either of these aircraft had insufficient spacing or was issued a go-around, however, this flight was captured in the daily PDARS go-around report for EWR. This data mining tool could be used to passively identify instances of inadequate spacing for post CRDA/Converging Runway Operations (CRO) analysis.

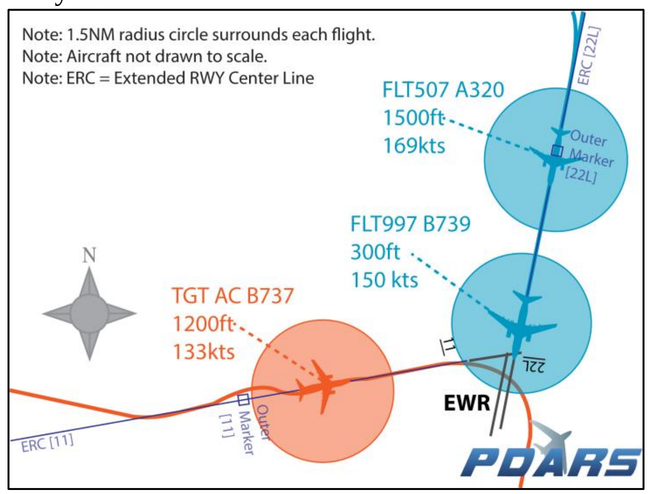

Figure 10 Converging Approach Go-Around Scenario.

DEN Scenario 10

Description: An Airbus A319 (TGT AC) on approach to DEN intercepts ERC for Runway 26 just inside the OM. FLT709 is on final to the same runway (3 NM ahead of TGT AC), but at a ground speed of 42 knots less than the TGT AC. The TGT AC continues through ERC and momentarily parallels the localizer before making a sharp left turn in the direction of runway 25 at the OM (not shown). TGT AC proceeds to cross Runway 25 ERC and executes a visual approach to runway 35R as FLT899 is at the same time on final to 35L.

Explanation: Typically the switch to Runway 35R could not be performed if a normal volume of traffic were present on 35R. The TGT AC seemed to be on a VFR pattern, which is also unusual for a busy airport such as DEN. Upon completing the turn inbound to Runway 35R, the TGT AC is 1.5 NM from

FLT889 and 300 feet below. Figure 11 shows an aircraft making a runway change that is not a normal procedure at DEN. The tool flagged this anomaly due to contributions from the unusual change in longitude and latitude.

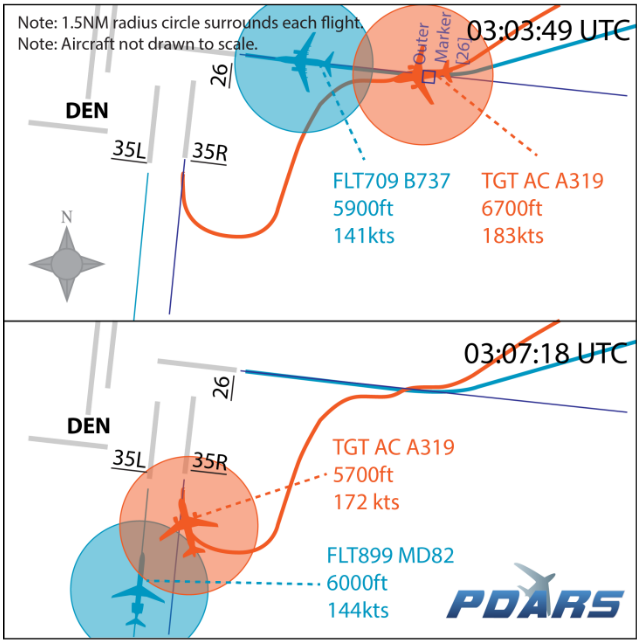

Figure 11 Abnormal Cross Runway Change Scenario.

Safety Review: The TGT AC ground speed at OM for Runway 26 was excessive and most likely would have led to an eventual overtake of preceding flight, ultimately resulting in go-around if runway switch was not made. The intercept angle for Runway 25 ERC at OM was 90 degrees making successful approach to runway 25 questionable. Visual approach to 35R was described as an "aggressive turn" by the tower controllers that reviewed this event. The event was close-in to a runway, at a high ground speed (172 knots) and a "belly-up" situation for TGT AC to FLT899 on the adjacent parallel approach.

DEN Scenario 11

Description: Figure 12 depicts a CR17 (TGT AC) intercepting the 17R ERC and a B190 (FLT01) intercepting the 17L ERC during period of triple approaches utilizing 16L, 17R, and 17L at DEN. At the specific time shown in Figure 12, the TGT AC has overshot the 17R ERC and does not have adequate separation from FLT01 by 700 feet laterally and 100 feet vertically.

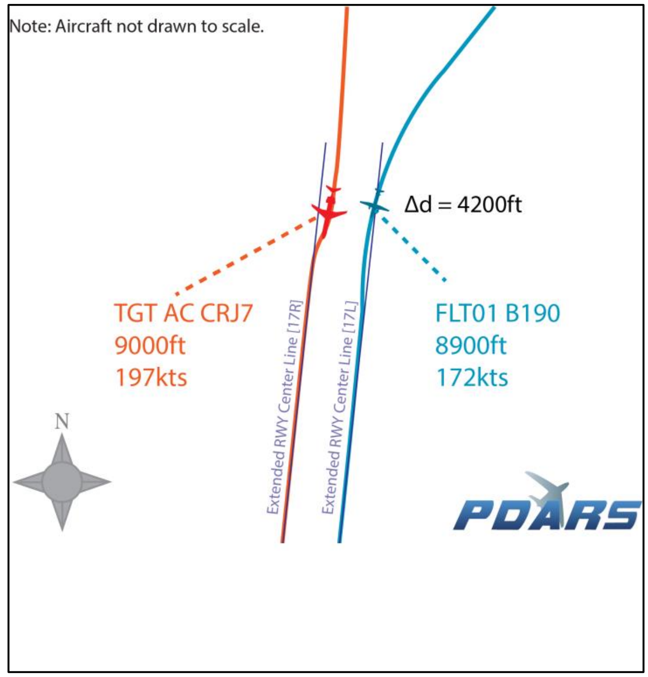

Figure 12 ERC Overshoot During Intercept of Parallel Runway Scenario.

Explanation: Five minutes prior to the overshoot event, the aircraft were separated by 7 miles and 1,000 feet. The TGT AC continued descending and approached the assigned ERC at 25 knots higher ground speed than FLT01. No voice recordings were available for this event. The tool flagged this anomaly due to its proximity to the nearest flight as well as contributions from the unusual shift in longitude.

Safety Review: Use of simultaneous parallel approaches requires that adequate spacing be maintained between the aircraft executing these approaches. In this case, the lateral and vertical distances were less than required. The pilot may have switched from RNAV path to approach mode (FMS) and then corrected localizer line up. The localizer signal strength would not have been an issue, since the usable distance on localizer at DEN is 30 NM and intercept was at 20 NM.

Contributing Factors

Figure 13 is a summary table that provides flight summary and characteristic categories for the 11 scenarios discussed above. For scenarios 1, 3, 6, 8 and 9 (the most common of the characteristic categories observed in the radar data), aircraft executed a go-around during their initial approach attempt. Similarly, in four scenarios (1, 4, 10, and 11) each of the flights exhibited characteristics of excessive ground speeds during ERC intercept or on final approach. Six scenarios experienced at least one occurrence of S-turns, too high on approach, and/or controller instructions being issued. Examples of other flight characteristics investigated by the subject matter experts as potential reasons for the flight anomalies that were identified by the algorithm and were observed more than once include: runway switching, excessive intercept angle, overshoot of the ERC, overtaking preceding aircraft, ground traffic interference, and separation issues. It is important to note that the algorithm can only identify contributing factors based on the features that it is given. For this study the algorithm is only analyzing high level trajectories. Additional features that can help distinguish finer grain safety risks such as those listed in Figure 13 could be included in future work to help automate the safety risk analysis.

	Scenario 1	Scenario 2	Scenario 3	Scenario 4	Scenario 5	Scenario 6	Scenario 7	Scenario 8	Scenario 9	Scenario 10	Scenario 11
Aircraft Type	CR17	E170/PA28	B772	B744/A320	B190	BE20	B739	E145	B737/A320	A319	B190/CR17
Season	Spring	Fall	Spring	Fall	Multiple	Spring	Summer	Summer	Summer	Winter	Fall
Time of Day	Noon	Morning	Mid-morning	Midnight	Midnight	Morning	Afternoon	Noon	Afternoon	Early Evening	Morning
Runway	22	En route	31R	04L/04R	Multiple	22L	22L	22L	11/22L	35R	17R/17L
Voice Available	Y		Y	Y	Y	Y		Y	Y	Y	
Runway Switch					Y					Y	
Go-around	Y		Y			Y		Y	Y		
360 inside OM					Y						
S-turn(s)					Y		Y		Y		
Too fast	Y			Y						Y	
Too high			Y	Y	Y			Y			
Intercept Angle excessive			Y	Y				Y		Y	
ERC Overshoot			Y	Y				Y			
Overtake	Y					Y				Y	
Not configured to Land	Y		Y								
Navigation system switch											Y
Ground/Converging Traffic					Y				Y		
Controller Decision/ clearance change					Y	Y	Y			Y	
Separation issue		Y									Y
Other traffic not insight						Y					
Wake Turbulence issue						Y					
Weather Factor			Y		Y	Y					

Figure 13 Safety Risk Summary Table for all 11 Flights.

Conclusions

This method has been demonstrated to be a valid and useful tool for identifying operationally significant safety events in approach trajectories landing at 4 major U.S. airports. The tool was able to discover events exhibiting flight characteristics indicative of safety-related concerns as well as incidents of unusual operational anomalies. The unusual operational anomalies identified included high energy approaches (speed and or altitude greater than normal for turn to final), unusual arrival routings, and unsuccessful CRDA spacing sequence of aircraft to intersecting runways. With the advancement of the daily reporting tool, in the demonstration PDARS environment, the method has been shown to provide an insightful glimpse into previously unmonitored potential safety risks in the NAS and has taken significant steps in becoming a tool that can be used to make informed decisions regarding safety risks automatically. The FAA has expressed interest in further testing and advancement of this tool on an FAA system at an additional TRACON facility. By utilizing this tool the FAA can help pre-identify safety risk trends to proactively mitigate these risks and help prevent future incidents or accidents. This tool addresses an aspect of safety monitoring that is not currently being leveraged and can help complement the current state-of-the-art by providing a more holistic approach to safety in an already tightly monitored system such as the NAS.

Future Work

With the implementation of the algorithm as a prototype daily reporting tool the anomalies identified every day will need to be validated to determine whether the algorithm is performing properly. In the future the flights identified by the algorithm should be evaluated by subject matter experts to determine their safety significance. Based on this feedback, changes to the algorithm may be needed. This might include finding a better kernel function or developing a way to distinguish operationally significant anomalies between purely statistically significant trajectories that pose no safety risks. NASA is currently looking into active learning methods that aid in improving the classification between these types of anomalies, which can re-rank the anomalies so the flights of interest are more interesting to safety analysts and relevant to critical safety risks in the airspace. Ultimately, successful identification of new areas of NAS safety risk from this approach can be integrated into future risk-based safety decision making methods being pursued by ANSPs including the FAA[11].

Acknowledgment

We would like to acknowledge the NASA Aeronautics Mission Directorate, Aviation Safety Program, System-Wide Safety and Assurance Technologies project for funding this work and the FAA PDARS Program Office for their approval to analyze the PDARS operational NAS data. We also would like to thank the subject matter experts: Raymond Mueller, Steve Wyloge, Glenn Hilgedick and on-site controllers for their insightful comments and perspective on the identified events. And finally we would also like to thank Peset Tan for his graphical expertise.

References

[1] Matthews, B., Srivastava, A., Schade, J., Schleicher, D., et al., "Discovery of Abnormal Flight Patterns in Flight Track Data," AIAA Paper No. DOI: 6.2013-4386 AIAA ATIO Conference, Los Angeles, CA, August 12-14, 2013.

[2] den Braven, W., and J. Schade, "Concept and Operation of the Performance Data Analysis and Reporting System," paper number 2003-01-2976, SAE Advances in Aviation Safety Conference (ACE), Montrèal, Canada, September 8-12, 2003.

[3] Statler, Irving C., "The Aviation System Monitoring and Modeling (ASMM) Project: A Documentation of its History and Accomplishments: 1999-2005," NASA/TP-2007-214556, NASA Ames Research Center, Moffett Field, California June 2007.

[4] FAA Report N. 7210.56C "Air Traffic Quality Assurance," prepared by the Air Traffic Evaluations and Investigations Staff, Chapter 7, August 15, 2002.

[5] Office of the Inspector General Report No. AV 2013 - 046, "FAA's Efforts to Track and Mitigate Air Traffic Losses of Separation are Limited by Data Collection and Implementation Challenges," February 27, 2013.

[6] Lishuai Li and R. John Hansman (2013) "Anomaly Detection in Airline Routine Operations Using Flight Data Recorder Data" MIT International Center for Air Transportation Report No. ICAT-2013-4

[7] Bhaduri, K., Zhu, Q., Oza, N., and A. Srivastava, "Fast and Flexible Multivariate Time Series Subsequence Search," IEEE International Conference on Data Mining, pp. 48-57. Sydney, Australia. 2010.

[8] Matthews, B., Das, S., Bhaduri, K., Das, K., Martin, R., Oza, N., et al., "Discovering Anomalous Aviation Safety Events using Scalable Data Mining Algorithms," AIAA Journal of Aerospace Computing, Information, and Communication. 2013.

[9] B. Matthews, S. Das, K. Bhaduri, K. Das, et al. "Discovering Anomalous Aviation Safety Events using Scalable Data Mining Algorithms", AIAA Journal of Aerospace Computing, Information, and Communication. 2013.

[10] Schölkopf, B., and A. Smola. "Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond," MIT Press, 2001.

[11] Borener, S., et al., "Design and development of an Integrated Safety Assessment Model for NextGen," paper presented at the International Annual Conference of the American Society for Engineering Management, 2012.